

SEMICONDUCTOR PACKAGE WITH A CONTROLLED IMPEDANCE BUS AND METHOD OF FORMING SAME

[0001] The present patent application is a continuation of, and claims priority to, U.S. Patent application no. 10/410,390 filed April 8, 2003 which is a continuation of U.S. Patent application no. 10/045,864 filed January 9, 2002 which is a divisional of U.S. Patent No. 6,404,660 dated June 11, 2002, all of which are incorporated herein by reference.

BRIEF DESCRIPTION OF THE INVENTION

[0002] This invention relates generally to the physical layout of semiconductor systems. More particularly, this invention relates to a high density planar semiconductor system with a controlled impedance co-planar interconnect channel.

BACKGROUND OF THE INVENTION

[0003] Semiconductor systems are implemented in a variety of configurations. One type of semiconductor system is a master-slave system in which a semiconductor-based master device controls a set of semiconductor-based slave devices. An example of a master-slave system is a memory system in which a master device memory controller coordinates the operation of a set of slave devices in the form of memory modules. By way of example, the invention is described in the context of a memory system, although the invention is equally applicable to other types of semiconductor systems.

[0004] As computer processors increase in speed, there is a growing burden being placed upon memory systems that provide data to computer processors. For example, video and three-dimensional image processing places a large burden on a computer memory subsystem.

[0005] One or more high frequency buses are typically employed to provide the required bandwidth in such systems. The higher the frequency of operation of the bus, the greater the requirement that the signals on the bus have high-fidelity and equal propagation times to the devices making up the subsystem. High-fidelity signals are signals having little or no ringing, and which have controlled and steady rising and falling edge rates.

[0006] Many obstacles are encountered in assuring the uniform arrival times of high-fidelity signals to devices on the bus. One issue is whether the bus is routed in a straight line

or routed with turns. Turns of the lines may not permit the construction of the bus lines in a way necessary to achieve uniform arrival times of high-fidelity signals to devices on the bus.

[0007] The assignee of the present invention has filed a patent application entitled "High Frequency Bus System", Serial Number 08/938,084, filed September 26, 1997, the contents of which are expressly incorporated herein. The foregoing patent application discloses a digital system 20 of the type shown in Figure 1. The system 20 includes a mother board 22, which supports a master device 24 and a set of slave modules 26A, 26B, and 26C. A bus 28 is routed in a horizontal and vertical manner to interconnect the master device 24 with the set of slave modules 26A, 26B, and 26C, as shown in Figure 1. The bus 28 is terminated in a resistor 30.

[0008] Figure 2 is a side view of one of the modules 26 of Figure 1. Module 26 has a set of slave devices 32A-32E mounted thereon. The slave devices 32 may be mounted on one side or both sides of the module. The module also includes a set of edge fingers 34 for coupling to the bus 28.

[0009] Figure 3 is a top view of one of the modules 26 of Figure 1. Module 26 has a set of slave devices 32A-32E mounted on it. Module leads 36 link the set of slave devices 32A-32E and thereby form a portion of the bus 28. Each slave device of Figure 3 is enclosed in its own package 33A-33E.

[0010] The structure of Figures 1-3 represents state-of-the-art packaging for master-slave systems, such as memory subsystems, which are operated with a memory controller (master) and a set of random access memories (slaves). Each slave device of Figure 2 is enclosed in its own package. Metal traces or module leads 36 are used between the packages.

[0011] Placing each slave device in its own package is relatively expensive. Furthermore, such an approach is relatively space-intensive. In addition, such an approach can result in substantive signal propagation delays between, for example, the first and last slave devices in a row of slave devices.

[0012] It would be highly desirable to improve the performance of semiconductor systems, such as master-slave systems in the form of memory systems. Such improvements could be exploited to support the increased information bandwidth of modern computers.

SUMMARY OF THE INVENTION

[0013] One embodiment of the invention defines an apparatus. The apparatus includes a first substrate having a set of semiconductor devices formed within it. The

apparatus also includes a second substrate. A third substrate has a data conductor coupled between first and second connections to the second substrate. The data conductor is coupled to the set of semiconductor devices at respective connection points.

[0014] Another embodiment of the invention is a semiconductor system with an interconnect channel with an input end to receive a set of input signals and an output end to route a set of output signals. A set of semiconductor devices are formed in a substrate, with each semiconductor device of the set of semiconductor devices being designed for independent operability and being electrically isolated within the substrate from adjacent semiconductor devices of the set of semiconductor devices. Each semiconductor device includes a set of input nodes and a set of output nodes formed on the surface of the substrate. The interconnect channel is positioned on the input nodes and the output nodes of each semiconductor device of the set of semiconductor devices. The set of input signals from the input end of the interconnect channel is applied to each semiconductor device of the set of semiconductor devices to produce the output signals at the output end of the interconnect channel.

[0015] The method of the invention includes the step of forming a set of semiconductor devices in a wafer. Each semiconductor device of the set of semiconductor devices is designed for independent operability and is electrically isolated within the substrate from adjacent semiconductor devices of the set of semiconductor devices. The forming step also includes the step of forming within each semiconductor device, a set of input nodes and a set of output nodes on the surface of the substrate. Faulty semiconductor devices and operable semiconductor devices are then identified within the set of semiconductor devices. A set of operable semiconductor devices that are adjacent to one another on the wafer are subsequently grouped. An interconnect channel is then applied to the set of operable semiconductor devices so as to electrically link input nodes and output nodes of each semiconductor device within the set of operable semiconductor devices.

[0016] The interconnect channel may be a set of metal traces positioned over the set of memory devices. Alternately, the interconnect channel may be formed on an interconnect substrate, such as a thin-film substrate, flexible tape, or a printed circuit board. The set of memory devices and the interconnect channel may be positioned in a single package.

[0017] The invention improves performance in semiconductor systems, such as master-slave memory systems. The improved performance increases information bandwidth of computer and computer subsystems. The invention produces high density systems with

reduced signal propagation times. The high density systems of the invention reduce packaging costs and improve thermal performance.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] For a better understanding of the invention, reference should be made to the following detailed description taken in conjunction with the accompanying drawings, in which:

[0019] FIGURE 1 illustrates a master-slave digital system constructed in accordance with the prior art.

[0020] FIGURE 2 illustrates a side view of a slave module from the master-slave digital system of Figure 1.

[0021] FIGURE 3 illustrates a top view of a slave module from the master-slave digital system of Figure 1.

[0022] FIGURE 4 illustrates a wafer with different sets of slave devices that may be packaged in accordance with the invention.

[0023] FIGURE 5 illustrates a wafer with uniform sets of slave devices that may be packaged in accordance with the invention.

[0024] FIGURES 6A and 6B illustrate high density planar semiconductor systems with co-planar interconnect channels in the form of metal traces in accordance with embodiments of the invention.

[0025] FIGURES 7A and 7B illustrate high density planar semiconductor systems of the invention incorporated into a system of the type shown in Figures 1-3.

[0026] FIGURE 8 illustrates buffering devices that may be used in accordance with an embodiment of the invention.

[0027] FIGURE 9 illustrates a high density planar semiconductor system of the invention connected to a motherboard in accordance with an embodiment of the invention.

[0028] FIGURE 10 illustrates a high density planar slave device system with a co-planar interconnect channel in the form of a thin-film in accordance with an embodiment of the invention.

[0029] FIGURE 11 illustrates a high density planar slave device system with a co-planar interconnect channel positioned in a package, in accordance with an embodiment of the invention.

[0030] FIGURE 12 illustrates a high density planar slave device system with a co-planar interconnect channel implemented with a thin film and bond balls, in accordance with an embodiment of the invention.

[0031] FIGURE 13 illustrates a high density planar slave device system with a co-planar interconnect channel implemented to isolate individual non-functional slave devices.

[0032] FIGURE 14 is an enlarged view of a portion of the device of Figure 13.

[0033] FIGURE 15 illustrates a wafer with different sets of slave devices separated by scribe lines.

[0034] FIGURE 16 illustrates processing steps for grouping, cutting, and packaging semiconductor systems in accordance with an embodiment of the invention.

[0035] FIGURE 17 illustrates a wafer with designated faulty dice identified in accordance with the process of Figure 16.

[0036] FIGURE 18 illustrates a wafer with grouped, functional dice identified in accordance with the process of Figure 16.

[0037] FIGURE 19 illustrates a wafer with interconnect channels positioned on grouped, functional dice identified in accordance with the process of Figure 16.

[0038] FIGURES 20A-20C illustrates the formation, from a wafer, of a set semiconductor systems with controlled impedance buses in accordance with an embodiment of the invention.

[0039] FIGURE 21 is a cross-sectional view of an embodiment of the invention.

[0040] FIGURE 22 is a cross-sectional view of a prior art system.

[0041] Like reference numerals refer to corresponding parts throughout the drawings.

DETAILED DESCRIPTION OF THE INVENTION

[0042] Figure 4 illustrates a wafer 40 used to fabricate a set of semiconductor devices 41. Each semiconductor device 41 is designed for independent operability and is electrically isolated within the wafer 40 from adjacent semiconductor devices 41. Thus, as used herein, the term “semiconductor device” or “integrated circuit” or “die” refers to a circuit with active components and connecting wires formed in a semiconductor material, which also forms a perimeter border without active components or connecting wires. The perimeter border without active components or connecting wires associated with each semiconductor device is referred to as a “scribe lane”. Typically, a semiconductor device is cut along a set of scribe lines within a corresponding set of scribe lanes to separate the semiconductor device from the

other semiconductor devices of a wafer. Figure 4 illustrates scribe lanes 39 surrounding each semiconductor device 41.

[0043] By way of example, each semiconductor device 41 may be a slave device, which may be a memory device, such as a DRAM or a SRAM. In the prior art, the wafer 40 is cut such that each slave device 41 is packaged separately. In accordance with the invention, groups 42 of slave devices 41 are positioned in a single package. As discussed below, the groups 42 of slave devices 41 are linked by an interconnect channel positioned on top of the slave devices. This configuration produces a high density semiconductor system with reduced signal propagation times. Furthermore, the configuration reduces packaging costs and improves thermal performance.

[0044] Figure 4 illustrates groups 42 of varying size. Figure 5 illustrates a wafer with groups 43 of slave devices of a uniform size, each including four slave devices 41 and a signal repeater 44. Observe that each slave device 41 and signal repeater 44 is surrounded by a set of scribe lanes 39.

[0045] Figure 6A illustrates a group 43 of slave devices 41 of the type shown in Figure 5. The group 43 comprises a set of slave devices 41A, 41B, 41C, 41D, and a signal repeater 44. The signal repeater 44 includes a set of transceivers 45. The set of slave devices 43 may be cut from a wafer 40, as shown in Figure 5. Figure 6A illustrates scribe lanes 39. Observe that the far left scribe lane, far right scribe lane, and horizontal scribe lanes in the figure have been used to cut the group 43 of slave devices 41. The interior vertical scribe lanes 39 have not been used to separate adjacent devices 41, but still exist between slave devices 41.

[0046] In accordance with the invention, an interconnect channel 54 is positioned over the set of devices 41A, 41B, 41C, 41D, and 44. In this embodiment of the invention, the interconnect channel 54 is implemented as a set of metal traces 56A, 56B, etc. The metal traces 56 may be formed on the wafer during the fabrication process. Alternately, they may be formed after wafer processing by using chemical vapor deposition or other well known semiconductor processing techniques.

[0047] The metal traces 56 forming the interconnect channel 54 constitute controlled-impedance transmission lines. The interconnect channel 54 is applied to a minimum of two slave devices, but may be routed over an entire wafer. Each slave device 41 includes a set of device bond pads or vias 58. The metal traces 56 electrically connect the device bond pads or vias 58 so that signals can be routed into the devices 41.

[0048] Figure 6B also illustrates a group 43 of slave devices 41A-4D and a repeater 44. The repeater 44 includes a set of transceivers 45 and a termination resistor R_T . The figure also illustrates vias or bond pads 58 connecting to the individual devices 41 and 44. Perimeter bond pads 58 are used to establish an electrical connection with a substrate bond pad 51 via bond wire 53. The substrate bond pad 51 is positioned on substrate 55.

[0049] Figure 7A is a top view of a slave module 26 of the type shown in Figure 3. However, in the embodiment of Figure 7A, a high density group of slave 43 is mounted on the module 26. The figure shows the bus 28 connected to interconnect channel 54. The impedance of the interconnect channel 54 is configured to substantially match the impedance of the bus 28. That is, the impedance of the interconnect channel 54 matches the impedance of the bus 28 within plus or minus 30%. In one embodiment of the invention, the interconnect channel has individual metal traces, each with a controlled impedance between 20 Ohms and 100 Ohms.

[0050] Observe in Figure 7A that the bus 28 is connected to the module 26 at a first device 41A and at the last device 41E. All other communication on the module 26 is via the interconnect channel 54. More particularly, Figure 7A illustrates that the bus 28 is routed from a bus segment 28A into the package 43 via a bus segment 28B. Next, the bus 28 is routed with conductor 54 to sequentially coupled slave devices 41A-41E. Finally, the bus 28 is routed out of the package 43 with a bus segment 28C, after which it is routed on the motherboard 26. This forms a transmission line environment with stubs 58 having a delay shorter than the edge rate of the transmitted signals.

[0051] The configuration of Figure 7A stands in contrast to the configuration of Figure 3, in which each slave device is in a separate package and module leads 36 formed on the module 26 are used to establish connections between slave devices. The prior art configuration of Figure 3 results in increased expense and prolonged propagation times compared to the system of the present invention, as shown in Figure 7A.

[0052] Figure 7B illustrates the bus 28 connected to an interconnect channel 54. In particular, the bus 28 is connected to the interconnect channel 54 via leads 59 of package 57. The interconnect channel 54 includes an interconnect substrate 72 (e.g., a thin-film device, flexible tape, printed circuit board, and the like) with vias 58 extending to individual devices 41A-41F. Device 41E includes a termination resistor R_T , while device 41F includes a set of transceivers 45. The value of the termination resistor R_T is selected such that its impedance substantially matches the impedance of the bus 28. That is, the value of the termination

resistor is selected such that its impedance is within plus or minus 30% of the impedance of the bus 28.

[0053] Figure 8 illustrates signal buffers that may be used in accordance with an embodiment of the invention. In particular, the figure illustrates a metal trace 56 of the interconnect channel 54 connected to the individual slave devices 41A, 41B, 41C, and 41D. Each slave devices includes an input signal buffer 47 and an output signal buffer 48. The signal repeater 44 includes a set of repeater buffers 49 to drive incoming and outgoing signals.

[0054] Figure 9 illustrates another embodiment of the invention. Figure 9 illustrates a set of slave devices 52A-52E. The slave devices 52A-52E are connected to the interconnect channel 54 through vias 58.

[0055] Figure 9 also illustrates connections between the semiconductor substrate 50 that includes the slave devices 52A-52E and a motherboard 62. A set of input/output bond pads 60 are used to establish electrical links with a motherboard 62 via bondwires 64.

[0056] Figure 10 illustrates an alternate embodiment of the invention in which the interconnect structure is implemented on an interconnect substrate 72. The interconnect substrate may be a thin-film device (e.g., polyimide), a flexible tape, a printed circuit board, or similar structure. The second substrate may be formed in a separate, optimized process and then be applied to the substrate 50.

[0057] In the embodiment of Figure 10, the slave devices have perimeter bond pads 70, while the interconnect substrate 72 has bond pads 74. Bond wires 76 are used to link the bond pads on the respective surfaces.

[0058] Figure 11 illustrates another embodiment of the invention. In Figure 11, the semiconductor substrate 50 is linked, via bond balls 84, to an interconnect substrate 80 in the form of a printed circuit board. Thus, in the embodiment of Figure 11, a flip-chip configuration is used. The advantages of a flip-chip configuration can therefore be exploited in connection with the invention. These advantages include reduced package inductance, compact packaging, and reduced manufacturing cost.

[0059] Figure 11 further illustrates that the substrate 50 and the printed circuit board 80 are positioned within a package 90. For example, the package may be a standard plastic pin grid array (PPGA) package. External pins 92 may be electrically connected to the printed circuit board 80 using conventional internal traces (not shown).

[0060] Figure 12 illustrates an alternate embodiment of the invention. In particular, Figure 12 illustrates a substrate 50 with device bond pads 70. An interconnect substrate 72, implemented as thin-film, flexible tape, or the like, is positioned on the substrate 50. A channel (not shown) is embedded in the interconnect substrate 72. The interconnect substrate 72 includes interconnect substrate bond pads 74. Bond wires 76 link the device bond pads 70 and the interconnect substrate bond pads 74. Further, the interconnect substrate 72 includes bond balls 84 for connection to a printed circuit board or other substrate. The embodiment of Figure 12 represents a hybrid package including bond wires and flip-chip functionality. This embodiment underscores how the features of the invention can be implemented into a variety of packaging configurations, thereby providing an engineer with many design options.

[0061] Figure 13 illustrates a technique of isolating a flawed slave device in a set of slave devices constructed in accordance with the invention. The figure shows a set of slave devices 52A-52F. Through standard testing of the slave devices it is determined that slave device 52D is flawed. Thus, it must be isolated from the remaining functional devices.

[0062] The circuit of Figure 13 includes a reset line 90, which is connected to each slave device 52A-52F. The figure also illustrates an interconnect channel 54, linking each slave device via a set of bond pads. Element 92 on each slave device is a single large bond pad symbolizing a set of small bond pads or vias.

[0063] Each slave device 52A-52F includes a scan-in pad 100 and a scan out pad 102. Bond wires 76 are used to connect a scan out pad 102 to a scan in pad 100 of an adjacent device. For example, as shown in Figure 13, a bond wire 76 is used to connect scan out pad 102A of slave device 52A to scan in pad 100B of slave device 52B.

[0064] In the case of a flawed slave device, such as slave device 52D in this example, a bond wire 76 is used to connect the scan in pad 100D to a ground pad 108. The slave device 52D is configured such that it is disabled when its scan in pad 100D is grounded. Notwithstanding the flawed and disabled slave device 52D, the remaining slave devices can operate by connecting a bond wire between the scan out pad 102C and the input bypass bond pad 104A. A lead 106 connects the input bypass bond pad 104A to the output bypass bond pad 104B. A bond wire 76 links the output bypass bond pad 104B to the scan in pad 100E of slave device 52E. This configuration accommodates a flaw in a slave device 52D in a set of slave devices 52A-52F.

[0065] Figure 14 is an enlarged view of a portion of the device of Figure 13. In particular, Figure 14 illustrates a portion of slave device 52D, including a scan in pad 100D

and a scan out pad 102D. An input buffer 120 is preferably connected to the scan in pad 100D, while an output buffer 122 is preferably connected to the scan out pad 102D. The figure also shows a lead 106 connected between an input bypass bond pad 104A and an output bypass bond pad 104B. It is appreciable from Figure 14 that the lead 106 is used to reduce the bond wire length when the device 52A is bypassed. In other words, in the absence of this lead 106, a relatively long bond wire would have to be used to span an inoperative device. By way of example, a relatively long bond wire 126 is shown in phantom in Figure 14. Such a bond wire is relatively difficult to implement, is susceptible to damage, and creates a relatively high impedance. The present invention avoids this problem through the use of the bypass lead 106, which is accessed through relatively short bond wires 76. The short bond wire 76A is connected between a scan out pad 102C and the output bypass bond pad 104A, while the short bond wire 76B is connected between the output bypass bond pad 104B and the scan in pad 100E.

[0066] Figure 15 illustrates a wafer 40 with different groups 42 of slave devices 41 separated by scribe lanes 39, which are shown as bold lines. Interior scribe lanes 38 also separate individual slave devices 41 within each group 42, as shown with dashed lines. The interior scribe lanes 38 are relatively narrow scribe lanes that will not be used for cutting the wafer 40. The interior scribe lanes 38 account for tolerance errors that occur during the step-and-repeat operation when masks are made. The interior scribe lanes 38 prevent adjacent die patterns from overlapping with each other.

[0067] An interconnect channel 54 connects individual slave devices 41 within a single group 42. Thus, for example, interconnect channel 54A connects individual slave devices 41A, 41B, 41C, and 41D of group 42A.

[0068] In the prior art, scribe lanes 39 are placed around each individual device 41. In contrast, with the present invention, standard scribe lanes 39 can be placed around a group 42 of devices 41, while relatively narrow interior scribe lanes 38 can be placed between devices 41 within a group, since the interior scribe lanes 38 will not be used for cutting. Space must be allocated on the wafer 40 for the scribe lanes 39. Observe that with the present invention the amount of space reserved for scribe lanes 39 can be reduced. This reduction in scribe lane space results in more chips per wafer, thus reducing cost and improving production capacity.

[0069] The novel structure of the invention facilitates new processes for testing, cutting, and packaging semiconductor devices. Figure 16 illustrates a testing, cutting, and

packaging process 140 utilized in accordance with an embodiment of the invention. The first step of the process 140 is to perform wafer level testing (step 150) of devices fabricated on the wafer. This step is performed in accordance with prior art techniques. The next processing step is to designate faulty dice (step 152). Figure 17 illustrates a wafer 170 with designated faulty dice 172 identified in accordance with step 152. In accordance with prior art techniques, the wafer 170 would be cut to form individual dice at this point. In contrast, the present invention groups functional dice (step 154) at this point.. The result of this process is shown in Figure 18, which illustrates the wafer 170 including grouped functional dice 180A, 180B, 180C, and 180D. In the present example, this process results in individual dice 182A, 182B, and 182C.

[0070] Interconnect is then applied to the grouped, functional dice (step 156). Figure 19 illustrates interconnect 54 attached to the grouped, functional dice 180. At this point, scribe lines can be defined within scribe lanes that will be cut (step 158).

[0071] The next processing step shown in Figure 16 is to cut the wafer (step 160). After the wafer is cut, grouped functional dice are packaged (step 162). Each packaged device is then tested (step 164). If an individual die is identified as non-functional at this point, a bypass is applied for the non-functional devices (step 166). This bypass operation is of the type described in connection with Figures 13-14. Observe that the bypass operation may be used to disable faulty devices identified during wafer level testing. Alternately, wafer level testing may be avoided altogether and faulty devices may be identified and isolated once they are in a package.

[0072] There are a number of inventive features associated with the process of Figure 16. First, observe that the grouping of functional dice is not performed until after faulty dice are identified. This allows non-defective sub-circuits to be grouped according to locations of defective sub-circuits. Next, note that the interconnect is applied to selected devices after appropriate grouping. Further, the process of Figure 16 flexibly defines scribe lines after the grouping of devices. The process of Figure 16 further includes the step of exploiting a bypass feature for non-functional devices. This step is advantageously performed after the semiconductor is placed within its package.

[0073] The advantages of the invention are more fully appreciated with reference to Figures 20A-20C. Figure 20A illustrates a wafer 40 with a set of semiconductor devices 41A-41F formed thereon. Scribe lanes 39 separate individual semiconductor devices 41. An interconnect substrate 72 is positioned over the wafer 40 and supports the interconnect

channel 54. Bond pads or vias 58 are used to establish electrical communication between the interconnect channel 54 and individual semiconductor devices 41.

[0074] Figure 20B illustrates a set of semiconductor devices 41A-41D cut from the wafer 40. The figure also illustrates the interconnect channel 54 positioned on the interconnect substrate 72. In this embodiment, the semiconductor devices 41A-41D are positioned on a motherboard 62. Electrical connections are established between the interconnect channel 54 and the motherboard 62 via bond wires 190.

[0075] Figure 20C illustrates a set of semiconductor devices 41E-41F cut from the wafer 40. The figure also illustrates the interconnect channel 54 positioned on the interconnect substrate 72. Once again, electrical connections are established between the interconnect channel 54 and the motherboard 62 via bond wires 190.

[0076] Figure 21 is a cross-sectional view of a substrate 50 (e.g., a first substrate) with a set of slave devices 52A-52N embedded therein. The signals from the slave devices 52A-52N are routed to a second substrate 62 (e.g., a printed circuit board or a motherboard) via an interconnect channel 54 (e.g., a third substrate or interconnect substrate). Observe in Figure 21 that the interconnect channel 54, for example in the form of a flexible tape, also operates to route signals between slave devices 52A-52N. Thus, the interconnect channel 54 operates as a multidrop bus for a single substrate 50 having a set of slave devices 52A-52N.

[0077] The configuration of Figure 21 stands in sharp contrast to a prior art implementation shown in Figure 22. Figure 22 illustrates a substrate 200A with a single slave device 202A embedded therein. Signals from the slave device 202A are routed to the motherboard 62 via package pins 206A-206N of the package 204A enclosing the substrate 200A. The motherboard 62 must support multiple packages 204A-204N. Each package 204 requires its own set of package pins 206A-206N for signal routing. Observe that the structure of Figure 22 requires signals between slave devices 202 (e.g., from 202A to 202N) to be routed over the motherboard 62. In contrast, with the system of Figure 21, these slave-to-slave signals may be routed internally with the interconnect channel 54, thereby bypassing the motherboard 62. This allows for greater processing speeds. In addition, it can be readily appreciated that the configuration of Figure 22 requires a large number of packages, whereas the system of Figure 21 can utilize a single package for all of the slave devices 52A-52N formed within the substrate 50.

[0078] Those skilled in the art will appreciate that various combinations of the embodiments of Figures 4-13 may be used in accordance with the invention. That is, various

techniques may be used to connect the interconnect structure to an external package or motherboard, and various techniques may be used to connect a package enclosing a set of slave devices to a motherboard or similar mounting structure.

[0079] In each embodiment of the invention, slave devices are combined to form a high density package. For example, the individually packaged slave devices 32A-32E of Figure 2 can be substituted with a single package 90, in accordance with the invention.

[0080] Those skilled in the art will recognize a number of advantages associated with the invention. First, the invention produces high density systems. The high density systems reduce signal propagation time. In addition, the high density systems reduce packaging costs, since a number of devices are positioned in a single package, instead of each device having its own package. Further, the structure of the invention improves thermal performance, since a larger semiconductor substrate, which is highly thermally conductive, is available to spread heat.

[0081] The traces forming the interconnect channel can be extremely narrow. This facilitates high density routing. The invention also reduces the number of external inputs and outputs since many signals will only be routed internally between semiconductors.

[0082] The invention facilitates various design optimizations. For example, if five slave devices are each implemented as a memory with a small associated controller, in many cases a single controller can be used for a set of slave devices. This sharing of resources is efficient, and it also improves manufacturing yield since simplified slave devices are easier to fabricate.

[0083] Note that memory systems typically employ an even number of slave devices. Grouping an odd number of slave devices allows for redundancy in case one slave device is found to be defective during manufacturing and assembly.

[0084] The invention improves performance in master-slave systems, such as memory systems. The improved performance can be exploited to increase the information bandwidth of computers and computer subsystems.

[0085] The foregoing description, for purposes of explanation, used specific nomenclature to provide a thorough understanding of the invention. However, it will be apparent to one skilled in the art that the specific details are not required in order to practice the invention. In other instances, well known circuits and devices are shown in block diagram form in order to avoid unnecessary distraction from the underlying invention. Thus, the foregoing descriptions of specific embodiments of the present invention are presented for

purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, obviously many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, the thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalents.